I

Deterrence: Surviving the Environment

2/Deterrence and Its Capability Requirements

I have already indicated that the first of the arms policy problems to be considered is that of deterrence, and it might be useful to introduce the subject with a brief discussion of inter-personal and inter-group influence, of which deterrence is but a particular form. In any competitive situation involving two or more players, each player will attempt to influence the behavior of the other(s) in one or both of two possible directions. He will attempt to *encourage* certain forms of behavior and *discourage* others. Unfortunately for men and nations, diplomacy has tended to focus primarily upon the *dis*couragement dimension of influence almost to the exclusion of *en*couragement.

It should be immediately evident that a strategy of deterrence is but a particular form of influencing by discouragement; more specifically, it is a strategy of threatened punishment or threatened denial, and proceeds from the premise that the anticipated but conditional destruction or denial of a player's values will discourage certain forms of behavior.[1] Conversely, it is assumed that encouragement can be offered by promised rewards, or even by the absence of punishment. Though strategies of encouragement play almost no role in military doctrine and will be essentially neglected in this part of the study, we shall return to this particular form of influence in Parts II and III.

Returning to the more negative aspect of inter-nation influence efforts—strategies of deterrence—we note that an attempt to inhibit, restrain, or deter is normally based on the techniques of intimidation by means of threatened punishment or denial. In military affairs, emphasis on punishment has been ubiquitous if not always successful. Its susceptibility to failure will be discussed in the next chapter, but here we shall assume that a strategy based on threatened punishment *will* succeed providing certain basic conditions are met; and in theory there is no reason to expect that the meeting of such conditions will produce anything but success.

THE DETERRENCE MODEL

Perhaps the most useful way of analyzing the decisional process by which anticipated outcomes influence behavior is that of the probability-utility model.[2] Whereas most such analytical efforts rely upon such gross and all-encompassing variables as risk or danger, this approach permits us to at least break these variables down into their components, and to begin by treating them one at a time. Thus we can say that a rational actor will choose among alternative courses of action on the basis of perceived or anticipated consequences. But these consequences or payoffs are not simply uni-dimensional; they have two complex dimensions. One dimension in the perceived payoff or outcome is the negative or positive value attached to it; in the lexicon, these are disutilities and utilities, respectively. The other, which is an equally subjective one, is the estimated probability that each conceivable outcome will occur.

To avoid any misunderstanding, let me discuss these two dimensions briefly. In examining the utility-disutility dimension, it must be emphasized that the positive or negative value which a decision-maker attaches to any anticipated outcome (or payoff) is a composite one; that is, it includes such utilities as the degree of damage he can do to the adversary's missiles and bombers or other strategic weapons, the number of casualties he can inflict on the adversary's population, the number of factories or transport facilities he can destroy, and the degree to which he can degrade the adversary's command and control network, to mention the more obvious ones. Since all of these diminish the adversary's capacity to strike back effectively at the attacker, they are positive values or "utilities" to the would-be attacker. Among the disutilities which the attacker must contemplate are the same sort of things in reverse, though he may weight his own population casualties and weapon losses differently than he does those he anticipates for the enemy. In a crude and intuitive sort of way, the would-be attacker is able to weigh, combine, and merge these various potential losses to the victim and to himself into what we call, respectively, utilities and disutilities.

In the simpler analyses of military strategy, one is led to believe that the would-be attacker compares the utilities (damage he can do to the victim) with the disutilities (damage the victim can do to him) and then makes the simple decision; that is, if expected utilities exceed expected disutilities, he considers it a satisfactory trade-off and decides to attack. But as has already been suggested, utility-disutility is only one of the two dimensions which will affect the choice between attacking and not attacking, for example. The other is subjective probability: in a sense, the odds which the decision-maker would give in betting on the outcome. Naturally, strategists attempt, by the use of operations research and related statistical techniques, to make their predictions as reliable, accurate, and objective as possible; but prediction regarding the military behavior of nations is far more difficult than prediction regarding the behavior of molecules, electronic circuits, or assembly lines. Hence these strategic predictions are still extremely rough and highly subjective.

---

[1] One of the few analyses of strategy to explicitly reflect any awareness of this range of alternative influence techniques is Glenn H. Snyder, *Deterrence by Denial and Punishment* (Princeton, N. J.: Center of International Studies, 1958). See also the same author's *Deterrence and Defense* (Princeton, N. J.: Princeton University Press, 1961), the most serious *empirical* effort to date to apply game theory to military strategy.

[2] This model is based on the probability-utility concepts used in mathematical approaches to psychology and decision-making, but makes no pretense of quantitative accuracy. See Robert M. Thrall *et al.*, *Decision Processes* (New York: John Wiley, 1954). For the application of a similar approach to military procurement problems, see Charles J. Hitch and Roland McKean, *The Economics of Defense in the Nuclear Age* (Cambridge: Harvard University Press, 1960), chap. x.

Having discussed these two dimensions as they affect the tendency of a decision-maker to attack or not to attack at any given moment, let us go a step further and ask: What are the kinds of outcome which he must first consider and then assign probabilities to? Obviously, these outcomes vary considerably depending upon the situation in which the decision-maker finds himself and the types of alternative actions he is considering. They would certainly be different if he was already fighting a limited war and was trying to decide for or against a particular tactical move, or if he was in a no-war context and was trying to choose between recognition and non-recognition of a revolutionary regime. But in today's strategic stand-off, which is the dominant concern of this study, the various alternative outcomes he must try to predict are of both a more decisive and a more dramatic nature. To put it bluntly, the strategists on both sides are constantly juggling subjective probabilities involving the near-obliteration of their nations.

More particularly, they are trying to predict what a successful attack would do to both sides, what an unsuccessful one would do, and what the relative probabilities of success and failure are. What, however, are the criteria of success and failure? By a successful surprise attack with strategic weapons, we mean one in which the great bulk of the victim's retaliatory weapons are destroyed (or withheld by choice) before they can be inflicted on the attacker. As the balance of this book will make clear, there can be no such thing as a totally successful attack upon the United States or the Soviet Union given the weapons technology of the period and the arms inventories possessed by each side. Thus "success" must be measured by subtracting (roughly speaking) the losses inflicted on the attacker by the damaged retaliator from the losses which the attacker has inflicted by getting in the first strike. One might also measure success by subtracting losses due to retaliation from the losses that would have been sustained had the attacker himself been the victim of a surprise attack. But for our purposes, a successful attack would be one in which "most" of the victim's strike-back forces were destroyed before they were ever launched; and an unsuccessful attack would be one in which only a "small" portion were so destroyed, and the attacker thus received a highly destructive retaliatory blow.[3]

Having identified these two dimensions, how does the decision-maker combine his subjective utilities and disutilities with his subjective probabilities of success and failure in order to reach periodic conclusions to these imperative questions? Obviously, no two individuals or decisional units will combine them in the same fashion, give the same weight to the variables that affect success or failure, or attach the same importance (consciously or otherwise) to each dimension. Not only will there be considerable variance in this calculation depending on the culture, the decisional process, the dominant sociopolitical values, and personal styles, but these factors may hardly be considered at all by either the careless or the irrational. Later on, we shall explore these deviations in some detail, but here we are trying to construct an analytical model by which we can understand and predict the behavior of a hypothetical rational decision-maker. Thus, for the moment, we shall assume that the choice between attack and no-attack for the potential aggressor will be a function of the product (again, in only the grossest terms) of subjective probability multiplied by the anticipated consequences of success or failure in utility-disutility terms.

---

[3] The use of such vague words as *most*, *small*, etc., should indicate both the inadequacy of contemporary criteria and the difficulty confronting the decision-makers. It should also suggest the latitude available to those who must attempt to weigh the trade-offs between people and missiles, and between hospitals and manned bombers. One might well ask, "How many lives would we give to destroy the forces which could destroy —— of our ICBM's?"

In graphic form, his decisional calculus would provide a rough index of the probability × utility product by producing a rectangle whose area represents that product. The utility he attaches to success is "measured" off on the right (plus) side of the horizontal axis and his disutility in the event of failure is marked off on the left (minus) side. And since we are considering only two alternatives, the subjective probability of success plus the subjective probability of failure must add up to 1 on the vertical axis. The higher one of these is, the lower the other becomes, with the total always one (certainty). Thus in the ideal case (from the deterrer's point of view) the alternative outcomes would appear more or less as indicated in Fig. 1 to the would-be attacker.

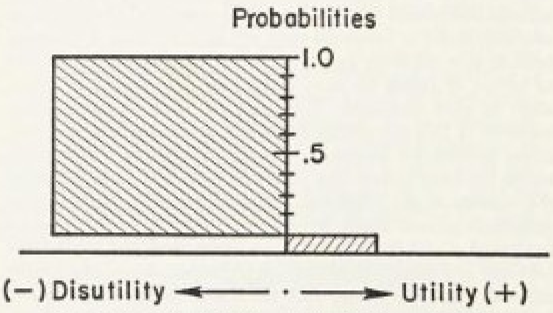

Fig. 1.—Deterree's estimated outcomes. Utility and probability of success are both low; disutility and probability of failure are both high. Rational deterree does not attack.

The situation portrayed here is a purely hypothetical one, with the disparity between the two probabilities and between the utility of success and the disutility of failure exaggerated for the sake of emphasis. If both sides could generate a set of estimated outcomes such as these in the minds of their opposite members in the cold war, the standoff would be a far more stable one than it is. But as the balance of Part I will emphasize, a multitude of factors makes this an utter impossibility.

Equally unlikely would be the opposite case in which the deterrer is able to generate not only a very low probability of failure for the would-be attacker, but a low disutility for him as well, even in the event that his surprise attack "failed." This, of course, would be an extremely unstable arrangement, with little incentive for restraint, and a powerful incentive to opt for the surprise attack; this is the "much to gain, little to lose" situation.

The more typical case in today's world is that reflected in Fig. 2, where the disparities are much less obvious. In this case, given the reciprocal offensive capabilities, the utility of success is valued by the deterree almost as much as he devalues the disutility of failure. Put in other terms, we can say that the positive value of destroying his enemy's capacity to destroy him is almost equal to the negative value of failing in that effort and having the enemy return a highly punitive retaliatory attack. But when the

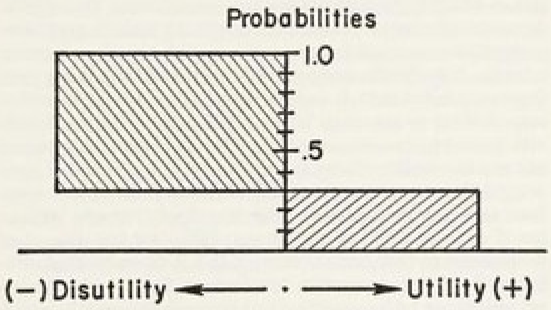

Fig. 2.—Deterree's estimated outcomes. Utility of success and disutility of failure are both high; probabilities are fairly equal. Deterree is ambivalent.

utility and disutility dimensions are more or less equal, the subjective probability of the two outcomes takes on an increasingly central role in the deterree's calculations. In this case, he estimates—in an intuitive fashion—that the probability of "success" is about .3; that is, he figures that the odds are about 7 to 3 against his being able to either destroy the bulk of his victim's retaliatory forces or induce such inhibitions that the victim will choose not to strike back at all. The result is that the area of the rectangle on the failure side is greater than that on the success side. Hopefully, the rational deterree is deterred. Note, too, that even if the utilities and disutilities remain the same, a moderate shift in his subjective probabilities might easily incline a calculating and amoral decision-maker toward the first-strike gamble.

So much for these illustrative cases here. Before turning to the factors which impinge upon these two dimensions in the rational deterree's calculus, let me reiterate briefly. All that we have been saying is that strategic deterrence is a special case of inter-group influence, but with the stakes much higher (some would say that they approach infinity) and the relative probabilities rather more difficult to calculate. If the relationship is asymmetrical, with both sides knowing who is in the role of possible attacker and who the possible victim, this model makes it reasonably clear what the latter needs to do in order to influence the former away from assuming that role—to deter him from strategic attack.[4] But if each sees itself as the deterrer, things are less clear. And if—as seems to be the actual case—each has an ambivalent self-image, the decisions are even less obvious. Thus, useful though our model may be for analytical and heuristic purposes, it cannot be taken too literally in the formulation of security policy.[5]

Now let us turn to the problem at hand and discuss the factors which operate to produce these estimates and expectations in the would-be attacker's decisional process. In the balance of this chapter, we shall focus upon the physical (with one exception) capabilities available to both the deterrer and the deterree, and discuss the ways in which such reciprocal capabilities act upon the two dimensions of the deterree's strategy calculus. In Chapter III, our emphasis will shift from the more or less objectively verifiable physical requirements which impinge primarily upon the utility-disutility dimension, to the more elusive and speculative variables which affect more markedly the second dimension in our deterree's calculations: the subjective probabilities he assigns to the various possible outcomes which he must contemplate.

This may not be the neatest or most obvious fashion of dealing with the probability-utility approach to deterrence, but it seems to offer several advantages. First, it emphasizes the distinction between those two discrete—but highly interdependent—dimensions of the requirements for deterrence. Second, it permits a nearly uninterrupted and systematic discussion of all the capability factors in one chapter, without numerous digressions. And third, it puts most of the dull and obvious material into a single chapter, permitting the knowledgeable reader to skip through it rapidly, merely noting the capability elements adduced, and the fashion in which they are arranged.

Capability requirements may be divided, for analytical purposes, into two parts: destruction and delivery. We shall examine first the destruction-level aspect, and make the temporary assumption that the retaliator is able to launch all of his bombers and/or missiles and deliver the nuclear payloads to, or over, the territory of the attacker. In the second part we shall deal with the delivery systems themselves and the problem of launching them on their retaliatory mission.

---

[4] We must always bear in mind that this simple model is useless if one or both sides derive some utility from gambling itself. See Halsey L. Royden *et al.*, "A Model for the Experimental Measurement of the Utility of Gambling," *Behavioral Science*, IV (1959), 11–18; and Ward Edwards, "The Reliability of Probability Preferences," *American Journal of Psychology*, LXVII (1954), 68–95.

[5] Another criticism may be that I have used a "sledgehammer to crack a pistachio nut," but the need for clarification seemed worth the effort. Such was the apology made by Richard C. Snyder *et al.*, *Decision-Making as an Approach to International Politics* (Princeton, N. J.: Foreign Policy Analysis Series, 1954), who also felt called upon to justify the use of an apparently overelaborate model.

Capability Requirements: Destruction Levels[6]

There would appear to be four major elements influencing the deterrent effectiveness of destruction levels: nuclear stockpiles available to the deterrer (potential retaliator), destruction levels acceptable to the deteree (potential attacker), and the passive and active defenses available to the attacker as he braces for the retaliatory blow. Let us discuss each in turn.

Retaliator's Nuclear Stockpiles

Perhaps the most dramatic and obvious factor is the total destructive force available to the deterrer. We shall start with this seemingly overwhelming quantity and then go on to examine the elements that begin to erode its impressiveness. As of the early 1960's, it is estimated that in nuclear bombs and warheads, ranging in size from low kiloton to multimegaton, the United States (with or without Britain and France) possesses a stockpile whose total yield could thoroughly obliterate, several times over, all relevant and vulnerable targets in the Soviet Union or its allies.[7] Such a destruction capability has often been referred to as "over-kill," but as we shall see in a moment, this may not always be an accurate characterization.

---

[6] This analysis of capability requirements is heavily indebted to the germinal studies by William W. Kaufmann, *The Requirements of Deterrence* (Princeton, N. J.: Center of International Studies, 1954); and Morton A. Kaplan, "The Calculus of Nuclear Deterrence," *World Politics*, XI (October, 1958), 20–43.

[7] The warheads used against Hiroshima and Nagasaki were each of about twenty kilotons yield: that is, they were equivalent in destructiveness to the effect of 20,000 tons of conventional high explosive. The term "megaton" is used to express millions of tons of high-explosive equivalent. A detailed analysis of the destructive power of a one-megaton warhead may be found in the testimony of Frank Shelton in *Joint Committee on Atomic Energy* (Special Subcommittee on Radiation), *Hearings . . .* , June 22, 1959, pp. 15–42. See also Samuel Glasstone (ed.), *The Effects of Nuclear Weapons* (Washington, D. C.: Atomic Energy Commission, 1957). One recent estimate made by Congressman Alger of the House Ways and Means Committee, sets the total stockpile at about 35,000 megatons, almost half of which is available for immediate retaliation (*New York Times*, November 11, 1961).

Attacker's Level of Acceptable Destruction

The second variable in our destruction-level equation is the degree of destructiveness which the potential aggressor is willing to accept in exchange for whatever destruction he is able to carry out.[8] This quantity is, of course, a function of more than just an educated guess as to how many casualties could be suffered without the total collapse of his nation. If he has a large population, much of which is not economically or militarily essential, he may be willing to accept personnel losses up to perhaps 75 per cent. Or if he has only just begun an industrialization program and he offers very few such targets, he may well find obliteration of his transport facilities a tolerable price to pay for certain types of "gain." Another factor in the acceptability of retaliatory destruction is the nature of the national and elite value systems. In societies long-accustomed to recurrent decimation by disease, hunger, revolution, guerilla war, or government terrorism, wholesale human slaughter may not be viewed by either the public or the elite as morally undesirable.[9] Some societies may see periodic population losses as inherent in the cosmic order, and some may even see them as a desirable way of eliminating the "unfit" or the "surplus." A final factor in the acceptability calculus is the political system. If it is of a highly authoritarian and centralized nature and the population is well controlled, the elites may in effect compel even a non-revolutionary or well-industrialized nation to accept astronomical human and material casualties. Thus it seems clear that tolerable destruction is a function of many factors, and that the acceptable level will range widely from one nation to another or within one nation over a period of time.[10]

---

[8] As I apologized earlier for the use of obvious anthropomorphizations, I must here express the hope that the reader will not object to my personification of states. The use of *he* saves many words and eliminates many distracting qualifications. It is understood throughout that the writer has in mind a rather complex decision-making apparatus at or near the apex of the various national governments.

[9] Raising the level of destruction that one finds tolerable produces one result that has been largely ignored: As side A's elite gradually comes to accept (publicly) an increased estimated loss, side B must, if he sees himself as the deterrer, raise his destructive capability so as to represent a sufficiently menacing threat of reprisal to deter A.

Attacker's Passive Defense

A third variable in the physical-capability equation revolves about the aggressor's capacity to reduce his losses by passive means. If he institutes a comprehensive civil defense, evacuation, and underground production program, the total damage a given nuclear megatonnage will do is appreciably diminished. There is, of course, considerable debate as to how effective such programs can be in reducing human and material losses. It has been argued that the warning times will soon be so short as to preclude evacuation of cities or retreat into shelters between the moment of warning and the moment of warhead impact, even for the side which strikes first. Roughly speaking, the intercontinental ballistic missile can travel between most Soviet launch sites and American targets (or vice versa) in slightly under thirty minutes. It is unlikely that the radar tracking stations now in operation would detect these missiles until they were about ten minutes from their destination, and the logistic problems of moving even a few hundred people out of a building and into a near-by shelter in so short a time seem insuperable to most critics of civil defense. Moreover, the expense and effort of building shelters capable of withstanding the percussion impact of multimegaton air or ground bursts are usually seen as prohibitive on both sides. In addition, there is the question of remaining alive within a shelter for the weeks that must elapse before it is comparatively safe to emerge. Moreover, the wholesale moving of industry underground or of dispersing it is seen to be so costly and dysfunctional that it is rejected by many. Despite these criticisms, however, there is no doubt that a passive defense system against nuclear attack would

---

[10] For some optimistic calculations concerning the amount of destruction that the United States could withstand and still "continue" as a society, see Herman Kahn, *On Thermonuclear War* (Princeton, N. J.: Princeton University Press, 1960), pp. 74–94. One of my criticisms of this study is that it contributes to the rising level of acceptable destruction mentioned in the previous note.

reduce somewhat the number of non-military casualties and the amount of destruction sustained.[11]

Attacker's Active Defense

The fourth major variable involved in our calculation of the destructive levels needed to deter is that of active defense. The aggressor may reduce the impact of the retaliatory blow he receives not only by protecting or dispersing his people and plants (as well as his military installations) but also by actually throwing up an umbrella which reduces the quality of megatonnage which reaches him. At this writing, with both sides still relying heavily (but to a rapidly declining degree) upon manned bombers to deliver the parcels of nuclear retribution, the number of such bombers shot down en route or over their targets by fighter planes or antiaircraft fire could reduce significantly the destruction the aggressor must accept. Even at extreme altitudes and speeds (50,000 feet, and 1,200 miles per hour for the B-58 medium jet bomber), the problem of interception is not great, and most estimates anticipate that the Soviet could probably destroy approximately one-half of any retaliatory bomber fleet directed toward its industrial or population centers, and perhaps more if the attack was concentrated on Soviet military installations.[12] The survival figure is, of course, a function of such elements as the number and type of bombers employed, their flight pattern, and the targets' dispersal and the readiness of the defenders. Furthermore, such technological elements as our use of the stand-off air-

---

[11] For a discussion of the requirements of a relatively effective civil defense program under present technological conditions, see Kahn, *op. cit.*, appendixes iii and iv. This program is based on Kahn *et al.*, "Some Specific Suggestions for Obtaining Early Non-Military Defense Capabilities and Initiating Long-Range Programs" (Santa Monica: Rand Corporation Research Memorandum RM-2206-RC). In the next chapter I shall go into more of the details of civil defense requirements and shall summarize Soviet efforts in this field.

[12] Air Force Chief of Staff Thomas D. White argues, on the other hand, that as (and if) the B-70 long-range bomber becomes operational, "it makes all of his . . . air defenses against the air-breathing threat obsolete." House Committee on Appropriations, *Department of Defense Appropriations for 1961, Hearings . . .* , Part 2, p. 313.